

Amendments to the Specification:

Please replace the paragraph beginning at page 31, line 21 and ending on page 32, line 2, which starts with "After 610" with the following paragraph that is between the two sets of dash symbols "--":

--After 610, the query manager determines (at 615) whether the table manager returned any replacement sub-network for the set of indices it received at 610. If not, the process transitions to ~~660~~ 655, which will be explained below. Otherwise, the query manager selects (at 620) one replacement sub-network retrieved by the table manager at 610.--

Please replace the paragraph beginning at page 72, line 13 and ending on page 73, line 3, which starts with "After 2240" with the following paragraph that is between the two sets of dash symbols "--":

--After 2240, the process determines (at 2242) whether it has examined all the graphs generated at ~~2242~~ 2202. If it has not examined all the graphs, it transitions to 2228 to select an unexamined graph. Otherwise, it discards (at 2244) multiple definitions of the same network or nearly the same networks. This is done by deleting all but one out of each group of generated sub-networks that have (1) the same graph table index, and (2) the same sets of function indices (defined at 2240). Such duplicate networks may appear for example because of symmetries of the graph structure. Based on the list of network table indices and function indices, the process then completes (at 2246) the database tables 525. Specifically, the process first creates the network table 1510, then the secondary index table 1505, and then the pivot-index table 1500. As mentioned above, the network table is sorted in an order specified by its stored network indices, the secondary table is sorted in an order specified by its stored primary and secondary indices, and the primary table 1500 is sorted in an order specified by its stored primary indices. After ~~2244~~ 2246, the process ends.--